Remarks

The application was filed with claims 1-5. Claims 1-5 are rejected. All rejections are traversed. All inventions in claims 1-5 are commonly owned by assignee. Claims 6-12 are new.

There are three totally unrelated ways that computer vision systems have been used on images that include faces.

The first ***detects*** faces from a cluttered background. This is a relatively simple background subtraction problem.

The second ***identifies*** faces. This is a slightly harder comparison problem.

The third is the very difficult task of ***classifying*** faces by gender differences.

The applicant respectfully requests the Examiner not to confuse detecting, identifying, and classifying faces.

These distinct fields of processing face images should not be confused. Those of ordinary skill in the art have recognized and assumed that different strategies and techniques need to be used in each of these distinct fields. The details in these strategies and techniques in which these fields differ will be explained in greater detail below.

In paragraphs 2-3, the Examiner rejected claims 1, 4, and 5 under 35 U.S.C. 103(a) as being unpatentable over "Training Support Vector Machines: an Application for Face Detection," Osuna et al. (Osuna), and "Gender Classification of Human Faces Using Hybrid Classifier Systems," Gutta et al., (Gutta).

Osuna only describes a system for detecting human faces. Osuna ends where the invention begins.

Osuna is irrelevant. Osuna detects a face in an image. We already know that there is a face in the image. Osuna has a simple problem. He picks a highly structured face, we all have eyes, eyebrows, nose, and mouth, out of a highly textured and cluttered background.

Unfortunately, the invention has to deal with a much more difficult problem of distinguishing subtle features in identically structured faces to determine gender. We do not need Osuna, nor can we use Osuna.

The only thing that Osuna might reveal is support vector machines (SVMs). But that is no revelation. In fact, SVMs predate computer vision, a field of the early 1990's. SVMs have been known and available to computer vision researchers for almost 40 years. Yet now, all of the sudden, after reading the application and without citing a single references to support his view, the Examiner has made the conclusion, which hundreds of experts in the field have failed to see, that the use of SVM in gender classification is obvious.

The reason that SVMs have not been used is that they are an unlikely choice. Conventional wisdom would have it that SVMs are a bad choice for gender classification.

Osuna describes how a support vector machines (SVMs) can be used to determine if an image contains a face. Osuna recognizes that "it is important to remark that face detection, like most object detection problems, is a difficult task due to the significant pattern variations that are hard to parameterize analytically."

Osuna makes a point of recognizing the problems with the prior art techniques that used neural networks, geometrical constraints, density estimation, labeled graphs, and clustering. As Osuna points out, one must maximize detection while minimizing false positives, and details such as choosing the number of free parameters, see page 134 left column, see more on this subject below.

Osuna uses a labeled database of images that contain faces and other images that do not contain faces for training. Osuna applies masking, gradient correction, and equalization in a preprocessing step. Osuna uses images of landscapes, trees, rocks, and buildings for the non-faces because they have highly textured patterns that are quite distinctive from faces, see page 134 right column. They use anywhere up **100,000** support vectors to get reasonable results.

Claimed is a method for classifying faces according to gender. These are different and distinguishable tasks. Those of ordinary skill in the art would recognize that distinguishing a face from a tree, rock, building or landscape is relatively easy since these, as Osuna states, have "different textured patterns." When classifying faces according to gender, one is dealing with slight differences in images that have identical textures. This is a much more complicated and different problem.

The Osuna training images have landscapes, trees, rocks and buildings. The claimed invention uses training images that only include faces. Those of ordinary skill in the art would immediately recognize that the claimed training images are useless to the Osuna system, and vise versa. In fact, the claimed images would make the Osuna system fail. The claimed test image always includes a face of unknown gender. Osuna does not care about gender. In addition, the Osuna test images may not even contain a face.

In summary, Osuna does not describe:

1) a method for classifying faces according to gender
2) training images including images of male and female faces
3) a hyperplane separating male and female faces
4) a test image of a human face of unknown gender
5) classifying the gender of the test image,
6) scaling and warping images
7) masking scaled images
8) reducing resolution

these are the elements of the claimed invention, none of which are shown by Osuna.

As stated above, *face detection* and *gender classification* are two entirely different problems and demand different solutions

Face detection requires a *heterogeneous and skewed* population of training data, i.e., images of all sorts of objects with far many more non-face examples than face examples. Face detection prefers a linear discriminant.

In contrast, gender classification requires a *homogeneous and symmetric* population of training data, i.e., all training images are human faces and both gender classes are represented equally in the training set. Gender classification demands a non-linear discriminant.

Also, the goals of these two tasks are different.

Face detection aims to achieve a *lowest false accept rate for a given false reject rate* for asymmetric decisions.

In contrast, gender classification aims to achieve *lowest overall misclassification error* in a more *symmetric* problem.

For those that understand SVMs, the boundaries obtained by an SVM in these two problems are likewise very different, as attested to by the fact that

the Osuna face detector requires *hundreds of thousands* of SVs whereas the invented gender classifier requires only *several hundred.*

Therefore, it can no longer be argued that one skilled in the art reading Osuna's face-detector paper will see it to be an obvious and natural approach for a gender classification task. In fact, common sense would steer one away from SVMs.

Likewise, Gutta's gender classification approach uses an entirely different methodology with radial basis functions and decision trees. These are orthogonal to SVMs.

In the prior art, it was generally believed that to discern faint gender difference, greater image resolution would be required than for face-detection, see Gutta page 1356 section 5, where 4608 pixels are required for his gender classifier.

Surprisingly, this is yet another rather "obvious" assumption, that is proven wrong by the present invention, which operates successfully with mere "thumbnails" for gender classification as small as 12x21 (252), again defying conventional wisdom.

For those that understand SVMs, the insight behind the invention lies in the realization that SVMs, under the right circumstances and correct implementation, *can* be used for gender classification. If used as described by

the invention, and not as used by Osuna, they can model highly complex and highly convoluted decision boundaries.

Up to now, it was generally believed that SVMs were only applicable to in cases were the test data is *heterogeneous and skewed* and the decision boundary is mostly *linear.*

The invention shows how SVMs can be applied to *non-linear homogeneous and symmetric* data.

At this point, it may be useful to compare Figure 2 in Osuna and Figure 2 as in the invention, as shown below. Please note the highly separated populations and the straight hyperplanes in Osuna, and compare with the intertwined populations and convoluted gender-bender hyperplane according to the invention.

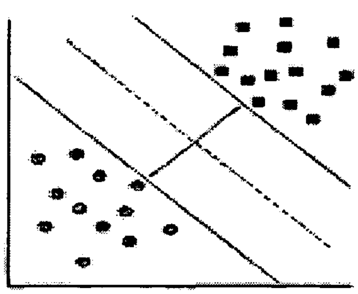

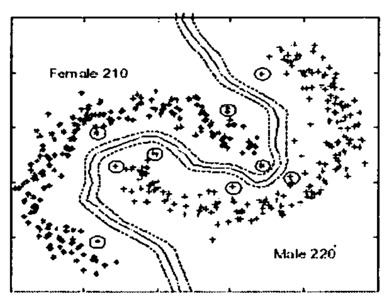

Prior Art SVM **Invention SVM**

Therefore, due to the lack of insight, up to know, prior art work in gender classification has concentrated on the use of extremely complex and sophisticated classifiers. A classical example of the best of them, exemplified

in the hybrid radial basis function/decision tree of Gutta, which has an at least four times larger test set than used by the invention, yet the invention outperforms Gutta by more than a factor of two.

The state-of-the-art, as reflected by the Gutta paper, was that the visual task of gender classification required complex and hybrid classifiers operating at high resolution and also including hair.

The invention attains better performance with a single non-linear SVM, operating on data with the lowest possible resolution and without the use of hair information.

With respect to claims 2-3, U.S. 5,710,833 by Moghaddam is also irrelevant. There, the task is face detection. In face detection, it is desired to isolate just that part of the face that does not include hair, and hair is considered part of the background.

Up to now, it has always been assumed that hair, or lack thereof, is an important part of gender classification in that hair styles of females, and males including those with mustaches, sideburns and beards, are quite different.

The invention makes the counter intuitive decision to exclude hair from the process, since it is too easy alter ones hair style and fool a classification system that includes hair as a factor. What is new is **not** *how* to remove hair, hair removal techniques are well know. What is new is to make the deliberate

decision to *minimize hair* during *gender classification*, a decision that, up to now, most gender classifiers have not dared to make.

All objections and rejections have been complied with, and Applicants respectfully submit that the application is now in condition for allowance. The Applicants urge the Examiner to contact the Applicants' attorney at phone and address indicated below if assistance is required to move the present application to allowance. Please charge any shortages in fees in connection with this filing to Deposit Account 50-0749.

Respectfully submitted for,
Mitsubishi Electric Research Laboratories

By: ______________________

Dirk Brinkman
Reg. No. 35,460
Attorney for Assignee

Mitsubishi Electric Research Laboratories
201 Broadway, 8th Floor
Cambridge, MA 02139
(617) 621-7573